

Efficient Decomposition of Bimatrix Games

Xiang Jiang & Arno Pauly
 Computer Laboratory
 University of Cambridge, United Kingdom
 Arno.Pauly@cl.cam.ac.uk

Exploiting the algebraic structure of the set of bimatrix games, a divide-and-conquer algorithm for finding Nash equilibria is proposed. The algorithm is fixed-parameter tractable with the size of the largest irreducible component of a game as parameter. An implementation of the algorithm is shown to yield a significant performance increase on inputs with small parameters.

1 Introduction

A bimatrix game is given by two matrices (A, B) of identical dimensions. The first player picks a row i , the second player independently picks a column j . As a consequence, the first player receives the payoff A_{ij} , the second player B_{ij} . Both player are allowed to randomize over their choices, and will strive to maximize their expected payoff. A Nash equilibrium is a pair of strategies, such that no player can improve her expected payoff by deviating unilaterally.

If the payoff matrices are given by natural numbers, then there always is a Nash equilibrium using only rational probabilities. The computational task to find a Nash equilibrium of a bimatrix game is complete for the complexity class PPAD [15, 5, 4]. PPAD is contained in FNP, and commonly believed to exceed FP. In particular, it is deemed unlikely that a polynomial-time algorithm for finding Nash equilibria exists.

The next-best algorithmic result to hope for could be a fixed-parameter tractable (fpt) algorithm [6, 9], that is an algorithm running in time $f(k)p(n)$ where n is the size of the game, p a polynomial and k a parameter. For such an algorithm to be useful, the assumption the parameter were usually small needs to be sustainable. The existence of fpt algorithms for finding Nash equilibria with various choices of parameters has been studied in [7, 11, 8].

In the present paper we demonstrate how *products* and *sums* of games – and their inverse operations – can be used to obtain a divide-and-conquer algorithm to find Nash equilibria. This algorithm is fpt, if the size of the largest component not further dividable is chosen as a parameter. *Products* of games were introduced in [17] as a means to classify the Weihrauch-degree [2, 3, 1, 12] of finding Nash equilibria for real-valued payoff matrices. *Sums* appear originally in the PhD thesis [18] of the second author; the algorithm we discuss was implemented in the Bachelor's thesis [13] of the first author.

2 Products and Sums of Games

Both products and sums admit an intuitive explanation: The product of two games corresponds to playing both games at the same time, while the sum involves playing *matching pennies* to determine which game to play, with one player being rewarded and the other one punished in the case of a failure to agree.

2.1 Products

In our definition of products, we let $[\cdot, \cdot] : \{1, \dots, n\} \times \{1, \dots, m\} \rightarrow \{1, \dots, nm\}$ denote the usual bijection $[i, j] = (i-1)n + j$. The relevant values of n, m will be clear from the context. We point out that $[\cdot, \cdot]$ is polynomial-time computable and polynomial-time invertible.

Definition 1. Given an $n_1 \times m_1$ bimatrix game (A^1, B^1) and an $n_2 \times m_2$ bimatrix game (A^2, B^2) , we define the $(n_1 n_2) \times (m_1 m_2)$ product game $(A^1, B^1) \times (A^2, B^2)$ as (A, B) with $A_{[i_1, i_2][j_1, j_2]} = A_{i_1 j_1}^1 + A_{i_2 j_2}^2$ and $B_{[i_1, i_2][j_1, j_2]} = B_{i_1 j_1}^1 + B_{i_2 j_2}^2$.

Theorem 2. If (x^k, y^k) is a Nash equilibrium of (A^k, B^k) for both $k \in \{0, 1\}$, then (x, y) is a Nash equilibrium of (A, B) , where $x_{[i_1, i_2]} = x_{i_1}^1 x_{i_2}^2$ and $y_{[j_1, j_2]} = y_{j_1}^1 y_{j_2}^2$.

Proof. We will prove that x is a best response to y , if x^k is a best response to y^k for both $k \in \{0, 1\}$, the remaining part is analogous. By applying the following equivalence transformation

$$\begin{aligned}
& \sum_{i=1}^{(n_1 n_2)} \sum_{j=1}^{(m_1 m_2)} x_i A_{i,j} y_j \\
&= \sum_{i_1=1}^{n_1} \sum_{i_2=1}^{n_2} \sum_{j_1=1}^{m_1} \sum_{j_2=1}^{m_2} x_{[i_1, i_2]} A_{[i_1, i_2][j_1, j_2]} y_{[j_1, j_2]} \\
&= \sum_{i_1=1}^{n_1} \sum_{i_2=1}^{n_2} \sum_{j_1=1}^{m_1} \sum_{j_2=1}^{m_2} x_{i_1}^1 x_{i_2}^2 (A_{i_1, j_1}^1 + A_{i_2, j_2}^2) y_{j_1}^1 y_{j_2}^2 \\
&= \left[\sum_{i_1=1}^{n_1} \sum_{j_1=1}^{m_1} x_{i_1} A_{i_1, j_1}^1 y_{j_1} \left(\sum_{i_2=1}^{n_2} x_{i_2} \right) \left(\sum_{j_2=1}^{m_2} y_{j_2} \right) \right] \\
&\quad + \left[\sum_{i_2=1}^{n_2} \sum_{j_2=1}^{m_2} x_{i_2} A_{i_2, j_2}^2 y_{j_2} \left(\sum_{i_1=1}^{n_1} x_{i_1} \right) \left(\sum_{j_1=1}^{m_1} y_{j_1} \right) \right] \\
&= \left[\sum_{i_1=1}^{n_1} \sum_{j_1=1}^{m_1} x_{i_1} A_{i_1, j_1}^1 y_{j_1} \right] + \left[\sum_{i_2=1}^{n_2} \sum_{j_2=1}^{m_2} x_{i_2} A_{i_2, j_2}^2 y_{j_2} \right]
\end{aligned}$$

on both sides of the best response condition

$$\sum_{i=1}^{n_1 n_2} \sum_{j=1}^{m_1 m_2} x_i A_{i,j} y_j \geq \sum_{i=1}^{n_1 n_2} \sum_{j=1}^{m_1 m_2} \hat{x}_i A_{i,j} y_j$$

one obtains the following form for the best response condition:

$$\begin{aligned}
& \left[\sum_{i_1=1}^{n_1} \sum_{j_1=1}^{m_1} x_{i_1} A_{i_1, j_1}^1 y_{j_1} \right] + \left[\sum_{i_2=1}^{n_2} \sum_{j_2=1}^{m_2} x_{i_2} A_{i_2, j_2}^2 y_{j_2} \right] \\
& \geq \\
& \left[\sum_{i_1=1}^{n_1} \sum_{j_1=1}^{m_1} \hat{x}_{i_1} A_{i_1, j_1}^1 y_{j_1} \right] + \left[\sum_{i_2=1}^{n_2} \sum_{j_2=1}^{m_2} \hat{x}_{i_2} A_{i_2, j_2}^2 y_{j_2} \right]
\end{aligned}$$

As this is just the sum of the best response conditions for the individual games (A^1, B^2) and (A^2, B^2) , the claim follows. \square

Theorem 3. If (x, y) is a Nash equilibrium of (A, B) , then (x^1, y^1) given by $x_i^1 = \sum_{l=1}^{n_2} x_{[i, l]}$ and

$y_j^1 = \sum_{l=1}^{m_2} y_{[j, l]}$ is a Nash equilibrium of (A^1, B^1) .

Proof. Again the proof uses contraposition. Assume w.l.o.g. that \hat{x}^1 is a better response against y^1 than x^1 , that is:

$$\sum_{i=1}^{n_1} \sum_{p=1}^{m_1} \hat{x}_i^1 A_{i,j}^1 y_j^1 > \sum_{i=1}^{n_1} \sum_{j=1}^{m_1} x_i^1 A_{i,j}^1 y_j^1$$

Add $\sum_{i=1}^{n_2} \sum_{j=1}^{m_2} x_i^2 A_{i,j}^2 y_j^2$ on both sides, and apply the reverse of the transformation used in the proof of Theorem 2. Then one obtains:

$$\sum_{i=1}^{n_1 n_2} \sum_{j=1}^{m_1 m_2} \hat{x}_i A_{i,j} y_j > \sum_{i=1}^{n_1 n_2} \sum_{j=1}^{m_1 m_2} x_i A_{i,j} y_j$$

with \hat{x} defined via $\hat{x}_{[i_1, i_2]} = \hat{x}_{i_1}^1 x_{i_2}^2$. This contradicts the assumption that x would be a best response against y , so (x, y) cannot be a Nash equilibrium. \square

2.2 Sums

The sum of games involves another parameter besides the two component games, which just is a number exceeding the absolute value of all the payoffs.

Definition 4. Given an $n_1 \times m_1$ bimatrix game (A^1, B^1) and an $n_2 \times m_2$ bimatrix game (A^2, B^2) , we define the $(n_1 + n_2) \times (m_1 + m_2)$ sum game $(A^1, B^1) + (A^2, B^2)$ via the constant $K > \max_{i,j} \{|A_{i,j}|, |B_{i,j}|\}$ as (A, B) with:

$$A_{i,j} = \begin{cases} A_{ij}^1 & \text{if } i \leq n_1, j \leq m_1 \\ A_{(i-n_1), (j-m_1)}^2 & \text{if } i > n_1, j > m_1 \\ K & \text{otherwise} \end{cases}$$

$$B_{i,j} = \begin{cases} B_{ij}^1 & \text{if } i \leq n_1, j \leq m_1 \\ B_{(i-n_1), (j-m_1)}^2 & \text{if } i > n_1, j > m_1 \\ -K & \text{otherwise} \end{cases}$$

Lemma 5. Let (x, y) be a Nash equilibrium of $(A^1, B^1) + (A^2, B^2)$. Then $0 < (\sum_{i=1}^{n_1} x_i) < 1$ and $0 < (\sum_{j=1}^{m_1} y_j) < 1$.

Proof. Assume (x, y) is a Nash equilibrium. The following circular reasoning demonstrates that any of the forbidden cases yields a contradiction.

1. If $0 = \sum_{i=1}^{n_1} x_i$, then also $0 = \sum_{j=1}^{m_1} y_j$.

If $y_j > 0$ for any $j \leq m_1$, then y' defined via $y'_j = 0$, $y'_{m_1+1} = y_{m_1+1} + y_j$ and $y'_l = y_l$ for $l \neq j, m_1 + 1$ is a better response against x than y : The payoff difference between y' and y is $y_j \left(\sum_{i_2=1}^{n_2} (B_{i_2, m_1+1}^2 + K) x_{n_1+i_2} \right)$, and by choice of K every $(B_{i_2, m_1+1}^2 + K)$ is positive.

2. If $0 = \sum_{j=1}^{m_1} y_j$, then $\sum_{i=1}^{n_1} x_i = 1$.

If $x_i > 0$ for any $i > n_1$, then x' defined via $x'_i = 0$, $x'_1 = x_1 + x_i$ and $x'_l = x_l$ for $l \neq 1, i$ is a better response against y than x : The payoff difference between x' and x is $x_i \left(\sum_{j_2=1}^{m_2} (-A_{i-n_1, j_2}^2 + K) y_{m_1+j_2} \right)$, and by choice of K every $(-A_{i-n_1, j_2}^2 + K)$ is positive.

3. If $\sum_{i=1}^{n_1} x_i = 1$, then $\sum_{j=1}^{m_1} y_j = 1$.

The proof proceeds as in 1. via symmetry.

4. If $\sum_{j=1}^{m_1} y_j = 1$, then $0 = \sum_{i=1}^{n_1} x_i$.

The proof proceeds as in 2. via symmetry.

□

Theorem 6. If (x, y) is a Nash equilibrium of $(A^1, B^1) + (A^2, B^2)$, then a Nash equilibrium (x^1, y^1) of (A^1, B^1) can be obtained as $x_i^1 = \frac{x_i}{\sum_{l=1}^{n_1} x_l}$ and $y_j^1 = \frac{y_j}{\sum_{l=1}^{m_1} y_l}$.

Proof. By Lemma 5, (x^1, y^1) is well-defined, and clearly a strategy profile. W.l.o.g. we assume that x^1 is not a best response against y^1 , and derive a contradiction. Let \bar{x}^1 be a better response against y^1 than x^1 . Define \bar{x} via $\bar{x}_i = \bar{x}_i^1 (\sum_{l=1}^{n_1} x_l)$ for $i \leq n_1$, and $\bar{x}_i = x_i$ otherwise. We claim that \bar{x} is a better response against y than x . The payoff difference between \bar{x} and x for the first player can readily be computed to be $(\sum_{l=1}^{n_1} x_l) \sum_{i=1}^{n_1} \sum_{j=1}^{m_1} A_{i,j} (\bar{x}_i^1 - x_i^1) y_{j2}$. Up to the positive factor $[(\sum_{l=1}^{n_1} x_l) (\sum_{k=1}^{m_1} y_k)]^{-1}$, this is equal to the payoff difference between \bar{x}^1 and x^1 , hence the former is positive iff the latter is. □

Theorem 7. Let (x^k, y^k) be a Nash equilibrium of (A^k, B^k) resulting in payoffs (P^k, Q^k) for both $k \in \{1, 2\}$. Then (x, y) is a Nash equilibrium of $(A^1, B^1) + (A^2, B^2)$, where $x_i = x_i^1 \frac{K-Q^2}{2K-Q^1-Q^2}$ for $i \leq n_1$, $x_i = x_{i-n_1}^2 \frac{K-Q^1}{2K-Q^1-Q^2}$ for $i > n_1$, $y_j = y_j^1 \frac{K-P^2}{2K-P^1-P^2}$ for $j \leq m_1$, $y_j = y_{j-m_1}^2 \frac{K-P^1}{2K-P^1-P^2}$ for $j > m_1$.

Proof. In the given situation, assume (x, y) were not a Nash equilibrium. W.l.o.g., let this be due to x not being a best response to y . This is equivalent to the existence of some i with $x_i > 0$, but the pure strategy i is not a best response to y . The latter means there is a better response k to y , i.e.:

Case $i \leq n_1, k \leq n_1$

$$\begin{aligned} & \left(\sum_{j_1=1}^{m_1} A_{kj_1}^1 y_{j_1}^1 \frac{K-P^2}{2K-P^1-P^2} \right) + \left(\sum_{j_2=1}^{m_2} K y_{j_2}^2 \frac{K-P^1}{2K-P^1-P^2} \right) \\ & > \left(\sum_{j_1=1}^{m_1} A_{ij_1}^1 y_{j_1}^1 \frac{K-P^2}{2K-P^1-P^2} \right) + \left(\sum_{j_2=1}^{m_2} K y_{j_2}^2 \frac{K-P^1}{2K-P^1-P^2} \right) \end{aligned}$$

Subtracting $(\sum_{j_2=1}^{m_2} K y_{j_2}^2 \frac{K-P^1}{2K-P^1-P^2})$ on both sides, then dividing by $\frac{K-P^2}{2K-P^1-P^2}$ shows that k is a better response against y^1 than i in the game (A^1, B^1) . But then $x_i^1 = 0$ follows, hence $x_i = 0$ in contradiction to the assumption.

Case $i \leq n_1, k > n_1$

$$\begin{aligned} & \left(\sum_{j_1=1}^{m_1} K y_{j_1}^1 \frac{K-P^2}{2K-P^1-P^2} \right) + \left(\sum_{j_2=1}^{m_2} A_{kj_2}^1 y_{j_2}^2 \frac{K-P^1}{2K-P^1-P^2} \right) \\ & > \left(\sum_{j_1=1}^{m_1} A_{ij_1}^1 y_{j_1}^1 \frac{K-P^2}{2K-P^1-P^2} \right) + \left(\sum_{j_2=1}^{m_2} K y_{j_2}^2 \frac{K-P^1}{2K-P^1-P^2} \right) \end{aligned}$$

As they are sums over stochastic vectors, we find $\left(\sum_{j_1=1}^{m_1} y_{j_1}^1\right) = \left(\sum_{j_2=1}^{m_2} y_{j_2}^2\right) = 1$. Moreover, $x_i \neq 0$ implies $x_i^1 \neq 0$, and this in turn implies $\left(\sum_{j_1=1}^{m_1} A_{ij_1}^1 y_{j_1}^1\right) = P^1$. Hence, after multiplying both sides by $(2K - P^1 - P^2)$ the previous inequality simplifies to:

$$K(K - P^2) + \left(\sum_{j_2=1}^{m_2} A_{kj_2}^1 y_{j_2}^2\right) (K - P^1) > P^1(K - P^2) + K(K - P^1)$$

This in turn can be simplified to $\left(\sum_{j_2=1}^{m_2} A_{kj_2}^1 y_{j_2}^2\right) > P^2$, which contradicts the assumption P^2 were the optimal payoff achievable by player 1 against y^2 .

Case $i > n_1, k \leq n_1$ Analogous to **Case** $i \leq n_1, k > n_1$

Case $i > n_1, k > n_1$ Analogous to **Case** $i \leq n_1, k \leq n_1$

□

3 The algorithm

Our basic algorithm proceeds as follows: To solve a game (A, B)

1. test whether (A, B) is the sum of (A^1, B^1) and (A^2, B^2) via some constant K . If yes, solve (A^1, B^1) and (A^2, B^2) and combine their Nash equilibria to an equilibrium of (A, B) via Theorem 7. If no,
2. test whether (A, B) is the product of (A^1, B^1) and (A^2, B^2) . If yes, solve (A^1, B^1) and (A^2, B^2) and combine their Nash equilibria to an equilibrium of (A, B) via Theorem 2. If no,
3. find a Nash equilibrium of (A, B) by some other means.

For some $n \times m$ game (A, B) let $S(A, B)$ denote its size, i.e. $S(A, B) = nm$, and let $\lambda(A, B)$ be the size of the largest game for which 3. in our algorithm is called. Let $f(k)$ be the time needed for the external algorithm called in 3. on a game of size k . Then the runtime of our algorithm is bounded by $O(S^3 f(\lambda))$, in particular, it is an *fpt*-algorithm:

Testing whether a game is a sum, and computing the components, if applicable, can be done in linear time. The sum of the sizes of the components is less than the size of the original game. Finally, combining Nash equilibria can be done in linear time, too.

Whether a game is a product of factors of a fixed size can also be tested in linear time. Testing the different possible factors yields quadratic time for this part. This already includes computing the components, and the product of the sizes of the factors is equal to size of the original game. Again, combining the Nash equilibria takes linear time.

As a slight modification of our algorithm, one can eliminate (iteratively) strictly dominated strategies at each stage of the algorithm. We recall that a strategy i of some player is called strictly dominated by some other strategy j , if against any strategy chosen by the opponent, i provides its player with a strictly better payoff than j . A strictly dominated strategy can never be used in a Nash equilibrium. It is easy to verify that a game decomposable as a sum never has any strictly dominated strategies, but may occur as the result of the elimination of such strategies. Hence, including an elimination step for each stage increasing the potential for decomposability. Elimination of strictly dominated strategies commutes with decomposition of

products, i.e. the reduced form of the product is the product of the reduced forms of the factors. The algorithm remains *fpt* if such a step is included. A detailed investigation of complexity issues regarding removal of dominated strategies can be found in [16].

4 Empirical evaluation

Only a small fraction of the bimatrix games of a given size and bounded integer payoffs will be decomposable by our techniques, this limiting the applicability of the algorithm in Section 3. However, to some extent we can expect patterns in the definitions of real-world game situation to increase the decomposability of the derived bimatrix games. For example, the structure of Poker-style games implies decomposability, as can be concluded from the considerations in [10].

To obtain a first impression whether using the decomposition algorithm is indeed beneficial, a collection of 100 random decomposable games was created. Each game has 95-105 strategies per player, and payoff values range from 0 to 50. The decomposability was ensured by creating a random tree representing the relevant decomposition structure first, using probabilities of 0.4 each for sum and product decomposition, and of 0.2 for an elimination of strictly dominated strategies step. The height of the trees was limited to 80, additionally vertices corresponding to games of size up to 6 were turned into leaves. At the leaves, the payoffs were chosen uniformly subject to the constraints derived from the structure and the overall constraint of payoff values being between 0 and 50. Finally, the corresponding bimatrix games were computed.

Both as a benchmark, and in order to compute Nash equilibria of the irreducible component games, the tool GAMBIT [14] was used. GAMBIT offers a variety of algorithm for computing Nash equilibria of bimatrix games, we used:

1. gambit-enmmixed: using extreme point enumeration
2. gambit-gnm: using a global Newton method approach
3. gambit-lcp: using linear complementarity
4. gambit-simpdiv: using simplicial subdivision

Figures 1.-4. show for each of the GAMBIT algorithms how many of our decomposable example games could be solved in some given time bound (per game, not total) using only the GAMBIT algorithm directly, or exploiting decomposition implemented in C++ first. Despite the fact that our decomposition algorithm was not optimized, it turned out that using decomposition almost all games could be solved in under 3 seconds, whereas even gambit-gnm as the fastest GAMBIT algorithm on the sample took 30 seconds for a similar feat. Thus, there is clear indication that on suitable data, exploiting the algebraic structure underlying the decomposition algorithm yields a significant increase in performance.

References

- [1] Vasco Brattka, Matthew de Brecht & Arno Pauly (2012): *Closed Choice and a Uniform Low Basis Theorem*. *Annals of Pure and Applied Logic* 163(8), pp. 968–1008, doi:10.1016/j.apal.2011.12.020.
- [2] Vasco Brattka & Guido Gherardi (2011): *Effective Choice and Boundedness Principles in Computable Analysis*. *Bulletin of Symbolic Logic* 1, pp. 73 – 117. ArXiv:0905.4685.
- [3] Vasco Brattka & Guido Gherardi (2011): *Weihrauch Degrees, Omniscience Principles and Weak Computability*. *Journal of Symbolic Logic* 76, pp. 143 – 176. ArXiv:0905.4679.




Figure 1: gambit-enummixed




Figure 2: gambit-gnm




Figure 3: gambit-lcp




Figure 4: gambit-simpdiv

- [4] Xi Chen & Xiaotie Deng (2005): *Settling the complexity of 2-player Nash-equilibrium*. Technical Report 134, Electronic Colloquium on Computational Complexity.
- [5] Constantinos Daskalakis, Paul Goldberg & Christos Papadimitriou (2006): *The Complexity of Computing a Nash Equilibrium*. In: *38th ACM Symposium on Theory of Computing*, pp. 71–78.
- [6] Rod Downey & Michael Fellows (1999): *Parameterized Complexity*. Springer.
- [7] Vladimir Estivill-Castro & Mahdi Parsa (2009): *Computing Nash equilibria Gets Harder – New Results Show Hardness Even for Parameterized Complexity*. In Rod Downey & Prabhuram Manyem, editors: *CATS 2009, CRPIT 94*.
- [8] Vladimir Estivill-Castro & Mahdi Parsa (2011): *Single Parameter FPT-Algorithms for Non-trivial Games*. In Costas Iliopoulos & William Smyth, editors: *Combinatorial Algorithms, Lecture Notes in Computer Science 6460*, Springer, pp. 121–124. Available at http://dx.doi.org/10.1007/978-3-642-19222-7_13.
- [9] Jörg Flum & Martin Grohe (2006): *Parameterized Complexity Theory*. Springer.
- [10] Andrew Gilpin, Javier Pena, Samid Hoda & Tuomas Sandholm (2007): *Gradient-based algorithms for finding Nash equilibria in extensive form games*. In: *Proceedings of the 18th Int Conf on Game Theory*.
- [11] Danny Hermelin, Chien-Chung Huang, Stefan Kratsch & Magnus Wahlström (2010): *Parameterized Two-Player Nash Equilibrium*. CoRR abs/1006.2063. Available at <http://arxiv.org/abs/1006.2063>.
- [12] Kojiro Higuchi & Arno Pauly (to appear): *The degree-structure of Weihrauch-reducibility*. *Logical Methods in Computer Science*. Available at <http://arxiv.org/abs/1101.0112>.
- [13] Xiang Jiang (2011): *Efficient Decomposition of Games*. Bachelor’s thesis, University of Cambridge.
- [14] Richard McKelvey, Andrew McLennan & Theodore Turocy (2010): *Gambit: Software Tools for Game Theory*. <http://www.gambit-project.org>. Version 0.2010.09.01.
- [15] Christos H. Papadimitriou (1994): *On the complexity of the parity argument and other inefficient proofs of existence*. *Journal of Computer and Systems Science* 48(3), pp. 498–532.
- [16] Arno Pauly (2009): *The Complexity of Iterated Strategy Elimination*. arXiv:0910.5107.
- [17] Arno Pauly (2010): *How Incomputable is Finding Nash Equilibria?* *Journal of Universal Computer Science* 16(18), pp. 2686–2710, doi:10.3217/jucs-016-18-2686.
- [18] Arno Pauly (2012): *Computable Metamathematics and its Application to Game Theory*. Ph.D. thesis, University of Cambridge.